WBS 3.10 - 1
An Overview of Distributed Computing in the
Engineering and Manufacturing Environment

An Overview of Distributed Computing in the Engineering and Manufacturing Environment

Prepared Under Contract No. NAS1-14700

WBS 3.10 Distributed Processing Research

By

Boeing Commercial Airplane Company

P.O. Box 3707

Seattle, Washington 98124

For

Langley Research Center

National Aeronautics and Space Administration

WBS 3.10-1

Prepared by:

David L Bernhardt Jr

D. L. Bernhardt, Jr.

DATE 3/24/83

Approved by:

Peter Dube

R. P. Dube, IPAD Software Manager

DATE 3/21/83

ABSTRACT

This paper is the result of the White Paper Scope milestone of the IPAD FY83 Task 3.10 Distributed Data Base Management Systems. It addresses four basic items. First, the objective of distributed computing in an engineering and manufacturing environment is discussed with respect to the characteristics of engineering and manufacturing; distributed computing; and layered networks. Next, the components of distributed processing are identified and each is discussed with an emphasis placed on standards and existing products. Work stations are discussed with discussion occurring on the types of work stations and their configuration possibilities. Finally, possible distributed procecessing architectures are presented which evolve from today's existing IPAD computing facility.

TABLE OF CONTENTS

TABLE OF CONTENTS

TABLE OF CONTENTS

LIST OF FIGURES

1.0 OBJECTIVE OF COMPUTING IN AN ENGINEERING AND MANUFACTURING ENVIRONMENT

Distributed computing in an engineering and manufacturing environment will be examined for the purpose of establishing the framework of the FY 83 design and analysis activity on Task 3.10 - Distributed Data Management. In this examination, the components of distributed processing will be addressed. This includes hardware, wide area and local area networks, distributed data base management system issues, and other related software. Next, the role of work stations in the distributed engineering environment will be examined. This examination will look at the broad types of work stations and configuration possibilities. A possible architecture for distributed computing will be presented which will include several phases of implementation. Using the distributed processing architecture as a basic foundation, a possible architecture for distributed data base management systems will be introduced. This will also include several phases of implementation.

However, before any of this can be explored, the engineering and manufacturing environment must be characterized so that a basic understanding of the role of automation in the environment is developed. Next, distributed computing in general must be inspected so that the motivations for it are understood. The differences between distributed processing and distributed data base management systems will be identified. Finally, the IPAD goal of a layered heterogenous computer network must be presented so that the context of the discussion can be maintained.

1.1 CHARACTERIZATION OF ENGINEERING AND MANUFACTURING ENVIRONMENT

A modern aerospace vehicle is a complex of sophisticated technical systems manufactured to the exacting standards required for safety, economy, and mission performance. The complexities of aerospace vehicles and processes of their design and manufacture have steadily increased with time. The development of more advanced vehicles is limited by the technology and methodology available to develop and manage the required analysis, design, and manufacturing processes.

1.1.1 Design and Analysis

Large design problems require a division of tasks. Individual designers are assigned a part of a task under the direction of a task leader. The designer performs his subtask by dividing it into several manageable jobs. Tasks are organized and reported by the task leader as part of a project group responsible for the design of a product.

The result of the design process is a single product entity in which all components must function harmoniously. Hence, extensive data communication, evaluation, and iteration are required to optimize the design. Orderly methods for introducing changes in the design are necessary to ensure a data (information) handling environment. This environment characteristically:

- o Requires extensive recordkeeping for data and configuration management;
- o Requires accurate and timely production and communication of data;
- o Has complex informal and formal communication paths;
- o Is continuous with respect to time, activities, and people.

These information functional characteristics remain unchanged even though the methods and computerized tools may change significantly.

As computing tools progressed, integrated systems were developed to relieve the constraints of communicating large volumes of data across program interfaces. These systems typically include several modules. Data can then be transferred across module interfaces by using compatible data formats and passing data through the computer directly.

These integrated systems can be configured to include several related disciplines, thus providing a medium for exchange of data and a common reference system for problem solution and theory development. By coupling these integrated systems with interactive management system, the computer-aided design and analysis functions can be coupled with numerical control software to provide a direct tie between engineering and manufacturing.

1.1.2 Manufacture

The manufacturing process starts with the release of design engineering information. This information is then used in various ways to generate the releases required to direct the manufacture of a product.

The primary functions involved in manufacturing are:

- o Production engineering
- o Fabrication
- o Accounting
- o Materiel
- o Status and scheduling

These functions produce a complex and large volume of data which has to be accessed by a large number of organizations.

Analysis of the present manufacturing process reveals that coordination across the entire manufacturing organization is limited. The various application computing systems do not communicate automatically, nor do they have a uniform description of data between them. This is because the activities in a department tend to operate in isolation. This leads to considerable redundancy of data stored when in fact it could be shared. The possibliity for inconsistencies and lack of integrity are enormous.

Since each department needs to view its data in a particular format, some means of providing multiple views are necessary to support the introduction of an integrated data management system.

Manufacturing data will continue to be spread between separate departments and activities and will continue to be processed on a wide range of computing facilities. Because of this, the need for distributed processing becomes very clear. The various options call for different degrees of distribution of both physical storage and control.

1.2 DISTRIBUTED COMPUTING

In this section, the general issues of distributed computing will be discussed, including the motivations for it and the distinction between distributed processing and distributed data base.

1.2.1 Motivations for Distributed Computing

There are many motivations for having a distributed computing environment. Furthermore, each situation is unique and has its own cost/benefit ratio which must be examined before a firm migrates toward a distributed computing environment. This section will address several motivations for distributed computing independent of any specific cost/benefit analysis. Although the motivations cited are felt to be important in the engineering and manufacturing environment, they are not all inclusive.

1.2.1.1 Local Control

Distributed computing tends to give organizations more control over the computing resource than does a centralized computing environment. Computing resources include: personnel, hardware, software, and their associated expenditures. Therefore, local control can be used to gain control of some of the computing costs. In addition to this, control is gained over the types of computing services provided to the organization including word processing, data capture, data management, and so forth. In a company where organizational groups tend to be autonomous in nature, local control of computing resources becomes highly desirable to the group's management. At the same time, this local control introduces some new problems which must be dealt with. These problems are mainly in the area of the integration of applications and the new obstacles to that process such as communication between heterogenous computers.

1.2.1.2 Integration of Turnkey Systems

Numerous CAD/Engineering work stations are available today and more are becoming available all of the time. These work stations are cost effective in certain environments. However, integrating them with other engineering and manufacturing applications increases their effectiveness, especially in the iterative design and analysis phases of the overall engineering task.

1.2.1.3 Response Time Requirements

Certain applications require fast response time if they are going to be effective. An example of this might be a graphics analysis package which displays airframe movement under varying load conditions. For this type of analysis to be effective, a high rate of communication between the CPU and display device is required. Seeing a line move at 2 minute intervals when it should be moving at 1 second intervals is enough of an impact to negate any potential benefit which may be gained from the analysis. Assuming a telephone type link, such slowness could be caused by the communications rate (e.g., 1200 Baud, 4800 Baud, etc.). Increasing line speed would improve performance but would also increase costs. Furthermore, it may not be possible to increase line speed to a high enough data rate to achieve sufficient terminal response. To achieve the necessary communications speed, a direct physical (no intervening modems and telephone lines) link between the CPU and the display device may be required. This, in many cases, may require moving either the computer closer to the users, or the users closer to the computer since most direct physical attachments are limited to short distances; less than one mile. In the case where neither alternative is possible, a small CPU in the same physical proximity of the user group may be introduced with the application and data moved to the local CPU. Having both the user and CPU in the same general area allows direct attachment of the display device to the CPU. As we now have a computer away from the central site, we are on the verge of distributed processing.

1.2.1.4 Reduced Down Time

Having the computing resource distributed among many machines as opposed to a single machine reduces the down time impact to a firm since any single machine which is down affects a much smaller community of users. However, in a distributed data environment, the backup and recovery problem is much greater.

1.2.1.5 Reduction of Communication Costs

The cost of communications may be reduced as a result of distributed processing. In the situation described in section 1.2.1.3 where a smaller computer is moved closer to the user, it is possible that multiple telecommunication lines, along with their corresponding modems and multiplexors, were replaced with fewer telecommunication links between the central site and the local site. This reduction in communication lines may result in a reduction of communication costs.

1.2.1.6 Sharing of Data

Distributed processing promotes the sharing of data in an environment where the portions of the integrated applications of a geographically dispersed engineering and manufacturing firm are also geographically dispersed. There are several points in the previous statement which bear directly upon the engineering and manufacturing environment. First, the environment is definitely geographically dispersed, sometimes as a matter of past events, other times as a matter of organization design. This geographic dispersion may be within the same firm or may result from a contractor/subcontractor arrangement or joint venture agreement. Second, the engineering and manufacturing community is moving increasingly toward the integration of design and manufacturing applications. The former is certainly an objective of the IPAD Project, which includes the industry emphasis in ITAB, and the latter is an emphasis of the Air Force ICAM Project as well as the emphasis of the Navy activities on the IPAD Project. Examples are all around us. All that we have to do is to examine the number physical locations in use by any one major aerospace firm.

1.2.2 Distributed Functions

The distributed function is a form of distributed computing where one of two conditions exist. In the first case, the function may be executed at more than one node in a computer network. In the second case, completion of a desired function may require work to be performed at various nodes of that network. The two cases are not mutually exclusive and most people tend to think of distributed processing as being a combination of both cases. An important aspect of distributed functions is that they must have knowledge of the network location of cooperative functions or data in order for them to accomplish the desired task. Data does not necessarily need to be distributed in a distributed function environment.

1.2.3 Distributed Data

Distribution of data may be accomplished using three basic methods: 1) distributed functions maintaining associated distributed data; 2) distributed data under data base management systems with local and centralized data bases; and finally, 3) distributed data base management systems. A necessary prerequisite for all methods is the ability to support distributed functions in a network.

In the distributed function method, applications control the distribution of the data, storing and retrieving it from known local or remote nodes. The distribution is therefore a function of the application design.

Using the second method, maintaining distributed data under multiple data base management systems, involves multiple copies of the data base. In this method, there is usually a centralized data base with local copies distributed to various nodes of the network. The local copies may be either complete copies of the data base or subsets of it with the subsets being complete units. In this scheme, either the summation of the local copies or centralized copy is designated as the "official" data base and the other is updated during the day. Then, during some low activity time period (e.g., 3-4 a.m.), the updated versions are reconciled against the "official" version or in other words synchronized.

The third method, a distributed data base management system maintains only one "official" version of the data base which is distributed over several nodes of the network. The data manager maintains the knowledge of the location of data and processes all user requests by making requests to other data managers on other nodes. Application programs can be oblivious to the fact that the data is distributed and its physical location in this environment. This methodology is by far the most complex and may be the most desireable long range solution. It is also the subject of much activity in the research community. Alan R. Hevner, editor of the December, 1982 issue of a quarterly bulletin of the IEEE Computer Society Technical Committee on Data Base Engineering, states in his introductory letter:

"Research in distributed database systems has enjoyed a short but very active history. In the past seven years, the research enphasis has been on establishing a theoretical basis for the important problems of distributed concurrency control, distributed query optimization, distributed system reliability, and design questions such as resource allocation on distributed systems. Numerous algorithms to solve these problems have been developed. As a result of this past work, research on distributed database systems has attained a level of maturity in which many of the important problems are now well understood. The challenge of current research becomes one of applying this knowledge to the design of experimental research activities."

1.3 LAYERED NETWORK

The ultimate objective to be achieved in an engineering and manufacturing environment is a layered network of computers. This network may consist of either homogenous or heterogeneous computers. Figure 1.3-1 depicts the concept. Layering of the network would make the best use of machine size, allowing small computers to perform simple tasks with large computers processing complex problems. Layering also works well in the "factory of the future" environment, allowing for specialized processors which monitor and control real time manufacturing processes. Heterogeneity would free the firm from being tied to a single vendor allowing the firm the latitude to make the most cost effective choice. Furthermore, ITAB has formally advised the IPAD Project to pursue a heterogeneous environment. In concept, a user (person or machine) would be unaware of the layers, and as more complex tasks were requested to be performed, the requested tasks would be performed on the appropriate machine in the network. Data would be ideally distributed throughout the network and managed by a distributed data base management system.

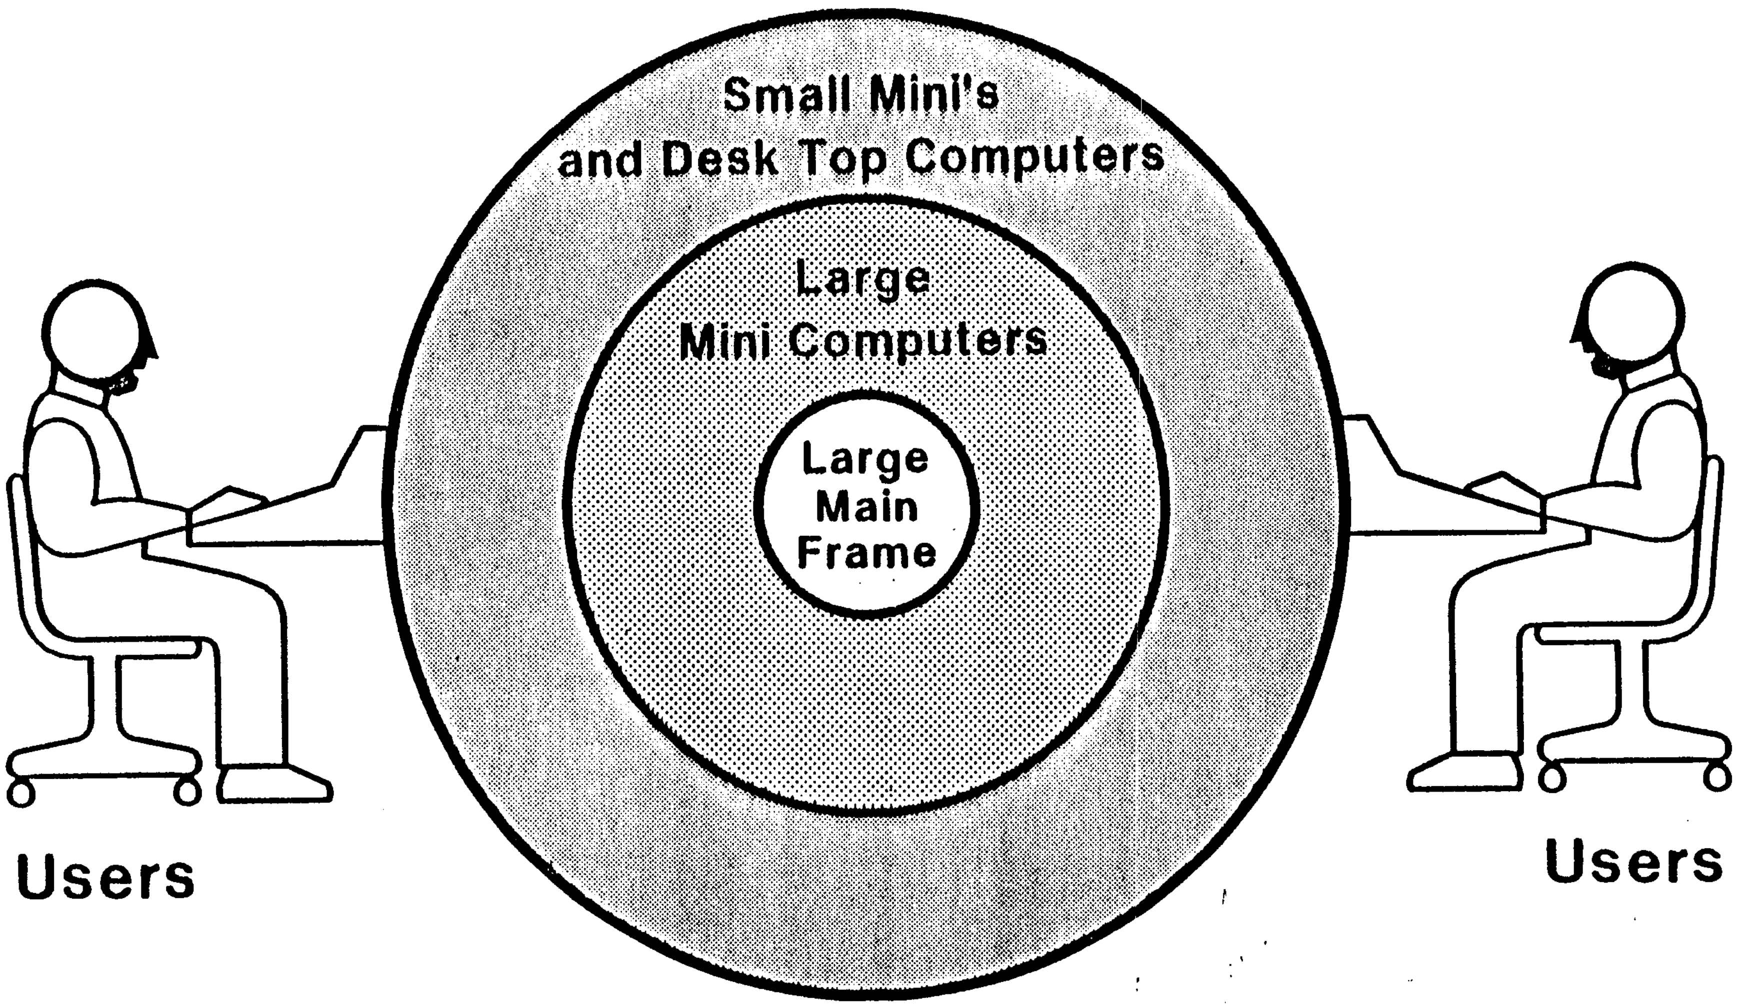

FIGURE 1.3-1 LAYERED COMPUTER NETWORK

2.0 DISTRIBUTED PROCESSING COMPONENTS IN A HETEROGENEOUS ENVIRONMENT

Distributed processing of any form is achieved through the integration of many components. These components include hardware; networks; software, including operating systems and teleprocessing monitors; and in some cases, distributed data base management systems. Figure 2.0-1 illustrates these components. The role of each of these components will be discussed relative to the engineering environment.

2.1 HARDWARE

Today's engineering and manufacturing environment contains, even in a small company, a broad spectrum of computing hardware. The hardware is provided by multiple vendors and includes general purpose computers with 8 to 64 bit words utilizing varying bit representations. These computers vary in size with a range from desk top micros to the largest available mainframe (such as IBM, CDC, UNIVAC, and CRAY processors). Also included in the spectrum of hardware are specialized turnkey systems such as design drafting systems, plotting systems, and word processors. Finally, there are the specialized small computers are used to monitor and control the manufacturing process. These specialized computers include NC and DNC machines along with their related cadre of hardware. This spectrum of machines is used in the process of designing and manufacturing a product.

2.2 NETWORKS

The ability to provide networking capability between computers requires the integration of many diverse items such as hardware, software, communication carrier mediums, and carrier disciplines. This becomes even more complex when networking heterogeneous computers because of differences such as character representation and word sizes. Furthermore, the market place has seen the introduction of and the beginning of the maturation of local area networks. Likewise, wide area networks are maturing and increasing in

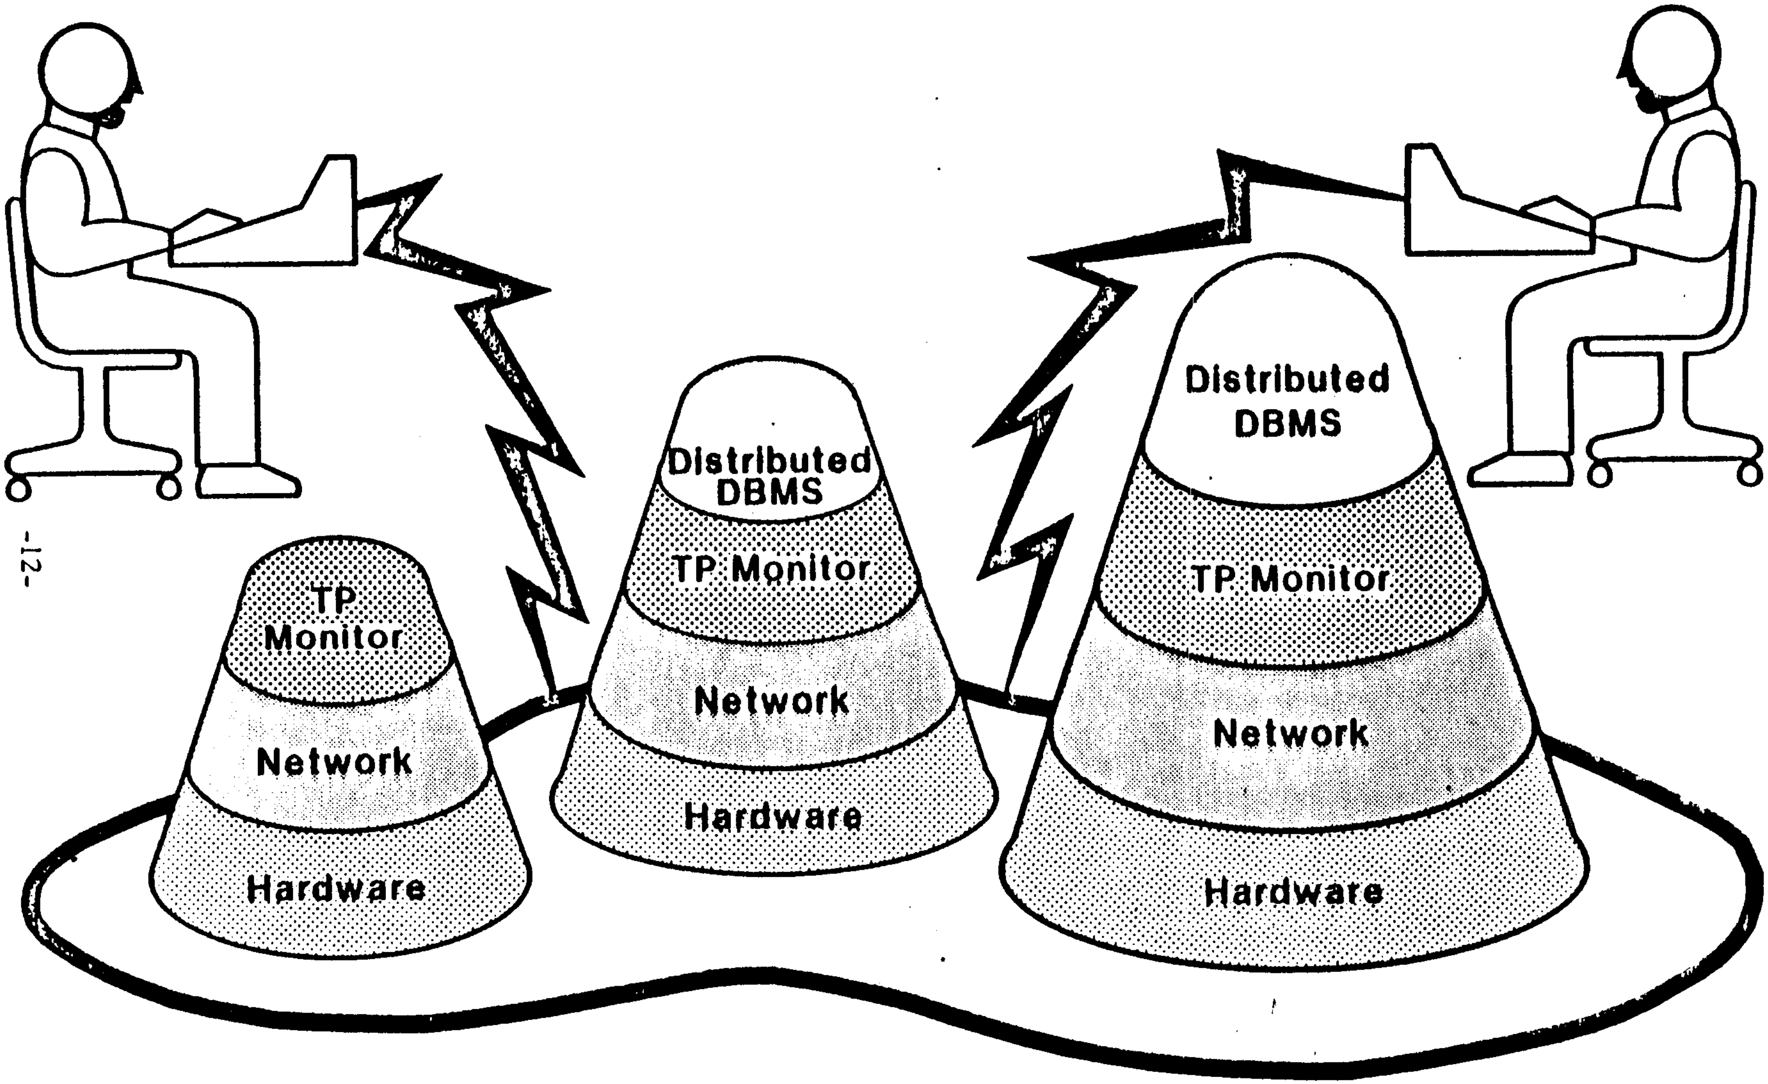

FIGURE 2.0-1 DISTRIBUTED PROCESSING COMPONENTS

capacity. Each of these two types of networks will be discussed, and where applicable, commercial products will be mentioned. However, prior to the discussion of the two types of networks, a certain level of understanding about network standards must be achieved.

2.2.1 Standards

The network field contains many standards. The three most well-known network standards are the International Standards Organization (ISO) reference model and the International Telegraph and Telephone Consultation Committee (CITT) X.21 and X.25 interface standards. The most recent development is the IEEE 802 standard. The ISO model provides for the separation of networking functions and a framework for the discussion of those functions. X.21 and X.25 provide standards for interfacing with wide area networks, and finally, IEEE 802 provides a standard for local area networks. Each will be discussed.

2.2.1.1 The ISO Reference Model

The ISO reference model provides a layered architecture for the implementation of network functionality. The seven levels of the model are:

Level 7	Application	
Level 6	Presentation	USER
Level 5	Session	
Level 4	Transport	
Level 3	Network	TRANSPORT
Level 2	Link	SERVICE
Level 1	Physical	

In general, Levels 1-4 provide transport services, while Levels 5-7 are users of the transport service. Each layer of the architecture provides a set of services to the layer above it using its own capabilities and the services of the layer beneath it. There is not necessarily direct one to one correspondence of the services of the lower level to the services of the upper level. For example, the application layer knows nothing about the physical layer, although the physical layer is required for communications to occur.

Level 7, the Application layer, is the user interface level. It provides support in the form of protocols to both applications and system functions. Examples of the system functions it supports are file transfer, remote batch processing, and distributed data base management systems. On the other hand, some examples of the applications it might support are electronic funds transfer and order entry.

Level 6, the Presentation layer, provides services for transforming the information being processed. Data encryption/decryption and terminal format changes are examples of the functions that this level performs. This level would also translate data formats and representations. Because of these translation services, the level becomes increasingly important in the heterogeneous environment when machine representations vary.

Level 5, the Session layer, supports the exchange of information between two processors. In this role, the level establishes, maintains, and terminates the sessions between processes. It may even initiate a process when required. This support also extends to the activities of checkpoint, recovery/restart, and the committment of a logical unit of data. Additionally, it buffers data and notifies processes of the receipt of data.

Level 4, the Transport layer, provides for the transport of data from one node of the network to other nodes in both simple and complex networks. It is a network independent interface to the network transport services.

Level 3, the Network layer, performs address to route resolution. It performs switching and routing determination in the network as opposed to the transfer of the data between the switching points.

Level 2, the Link layer, provides for the interchange of data between equipment. This is the actual carrier discipline involved such as Synchronous Data Link Control (SDLC) or Asynchronous.

Level 1, the Physical layer, is the physical, electrical, and other necessary items required to make the physical connections between two nodes of the network.

Using this architecture, it is possible to model a simple network or a complex one consisting of many nodes.

2.2.1.2 X.21

The CITT X.21 standard, adopted in 1976, provides specifications for a digital signaling interface using telephone lines. This standard specifies how terminals and computers are physically connected to and communicate with carrier equipment. The key to the standard is that it provides for an interface to digital communications rather than the current analog communications. Since digital communications were not in wide use in 1976, CITT also adopted an interim specification for analog communications known as the X.21 BIS. This specification is similar to the Electronics Industries Association (EIA) RS-232C specification. The RS-232C specification is generally used as a specification for electrical interconnection of data terminal equipment (DTE, e. g., terminals) and data circuit terminating equipment (DCE, e. g., modems). It is equivalent to the CITT V.24 electrical interface. As all of these standards involve the physical and electrical interfaces to telecommunication links; they correspond to the Physical Link layer of the ISO model.

2.2.1.3 X.25

The CITT X.25 standard provides a standard for packet-switched networks. As such, it covers the first 3 layers of the ISO model (Physical, Link, and Network) and in some cases, covers the Transport layer. The standard references both the X.21 standard and the X.21 BIS to cover the Physical link.

The Data Link level portion of the standard proposes a Link Access Protocol (LAPB) protocol which evolved from IBM's Synchronous Data Link Control (SDLC). Finally, the Network layer provides for data routing and virtual circuit management by defining a protocol, implemented with packets. Some examples of this protocol are: CALL REQUEST, CALL ACCEPTED, CLEAR REQUEST, and CLEAR CONFIRMATION.

2.2.1.4 IEEE 802

The IEEE 802 standard which has been proposed provides a standard for base band local area networks (see section 2.2.2.3 for a definition of base band). This standard appears to cover Level 2 and 3 in the ISO model, and closely parallels the Ethernet specification which was developed jointly by XEROX, Digital Equipment Co., and INTEL. Recently, 13 firms announced their endorsement of the standard.

2.2.1.5 Other Standards

The standards which have been described are not the only standards which exist. CITT has the following standards which support the X.25 standard. First, is the X.3 Packet Assembler Disassembler (PAD) standard for a piece of hardware which sits between the terminal and an X.25 network. The X.28 standard governs communications between the "dumb" terminal and the PAD, while the X.29 standard provides for communication between the PAD and the X.25 network. There is also in existence many other CITT standards. Likewise, IEEE has many standards in addition to the 802 and will most likely produce a standard for broad band local area networks (see section 2.2.2.3 for a definition of broad band).

2.2.2 Local Area Networks

Local area networks have undergone considerable evolution in the last few years. At present, there are over 49 offerings in the marketplace by more than 43 vendors with more appearing every day. There are industry rumors that major firms (such as IBM and AT & T) will soon enter the marketplace. The current offerings provide differing capabilities. Some are turnkey systems which provide word processing support, while others require software to be developed. Some local area networks provide services in all seven layers of the ISO model, while others provide services in levels one through three. In general, the local area networks have three things in common. First, as their name implies, they operate in a limited distance environment with a maximum of one to two kilometers distance between two nodes. Secondly, they provide relatively high transmission rates. These rates are generally between one and fifteen megabits per second, although some provide transfer rates up to 50 megabits per second. Finally, in their implementation in the user community, they are generally implemented within a single organization. These three items serve to provide a definition of what a local area network is.

2.2.2.1 Motivations for the Use of Local Area Networks

The motivations for use of local area networks are many. First, is the integration and exploitation of distributed functions and data which may be resident on a variety of mini and microcomputers located in an organization's area of responsibility. Secondly, is the sharing of peripherals among numerous mini and microcomputers, word processors, and the resulting reduction in costs. For instance, micros and word processors may share hard disks and letter quality printers, thereby reducing the total number of units required. Several micros might share a gateway communication link to a mainframe, and thus, eliminate the need for separate modems for each micro. Many more such examples exist. These reduced costs are offset by the increased cost for the local area network, but in many cases, a net reduction may occur. However, each situation must be examined for its own cost/benefit ratio. Finally, using local area networks rather than telephone communication links provides the using organization an additional element of control over communications and related costs.

2.2.2.2 The Physical Components of Local Area Networks

There are two basic components of Local Area Networks: the carrier and the hardware.

Local Area Networks have been implemented using five basic carriers. These five are: coaxial cable; dedicated twisted wire pairs; optic fibers; Private Branch Exchange (PBX); and finally, common carrier. Coaxial cable appears to be the most commonly used carrier in those products providing the higher transmission speeds, with twisted pairs used in products generally providing lower speeds. PBX is used by relatively few offerings but provides both voice and computer communications over telephone lines. Research work is being done on optic fiber technology, but this has yet to become commercially available. Common carriers are generally used to provide long distance gateways between local area networks and wide area networks or mainframes.

The hardware used in the creation of local area networks is as varied as are the types of local area networks. It does, however, consist of 2 basic devices. First, is the carrier tap; a device which allows the physical connection to a carrier. Secondly, there is the Network Interface Unit (NIU). A NIU generally has some intelligence in it (either software or firmware) and serves to: 1) interface with other devices; and 2) implement the network protocol. NIUs come in different varieties (i.e., hard disk, gateway, CPU, etc.). Some NIUs are programmable while others are not. Those NIUs supplied with turnkey systems are generally not programmable. NIUs may be free standing or boards which are plugged into existing hardware. The specific characteristics of a NIU are dependent on the local area network vendor and the device being interfaced.

2.2.2.3 Protocols and Products

There are many protocols used by local area networks. This will not attempt to delineate all of them but merely point out what is felt to be two of the more popular ones. These protocols appear to cover Levels two and three of the ISO model. Generally speaking, local area networks are divided into two types: broad band and base band. This refers to the band width of the signal sent over the carrier. All known broad band products use a coaxial cable carrier. (Strictly speaking, to be categorized as either base band or broad band requires a coaxial cable carrier.) The two major competitors in the marketplace, Ethernet and Wangnet, have taken different paths. Ethernet has chosen a base band approach while Wangnet uses a broad band approach. (Hyperchannel was not considered due to its specialized market segment.) Base band systems treat the carrier as a single high speed channel sending the signal out in digital form. Only one signal is allowed on the carrier at a time. Ethernet and Ethernet compatible systems use a Carrier Sense, Multiple Access, with Collision Detection (CSMA/CD) protocol for message passing. This is basically a time division, multi-plexing scheme which listens to the network, and when it is clear, sends a message. If a collision occurs (collisions are detected), the message is sent again at a random time interval. This protocol was developed by XEROX, Digital Equipment Corp., and INTEL and was recently proposed as the IEEE 802.3 standard for base band local area networks. This creates a somewhat confusing environment since Ethernet is both a generic name and a XEROX product name. On the other hand, it creates a marketplace with a host of compatible products.

Broad band systems separate the band width of the carrier using frequency division multi-plexing techniques to create many "bands". Using this technique, Wangnet has separated the total band width into four bands: the Wang band; the Utility band; the Dedicated Interconnect band; and the Switched Interconnect band. The Wang band is used to interconnect Wang computers and employs a CSMA/CD technique in that band alone. The utility band can carry color TV signals on seven channels, while the interconnect bands allow for dedicated and switched connection to computer devices through the use of special modems. The dedicated interconnect band provides up to 48 channels while the switched interconnect band provides up to 256 switched channels. Wangnet uses two coaxial cables for its carrier: one to listen; and one to send.

Neither of the two products is better than the other, as each represents a different view of the local area network philosophy. However, given the range of Ethernet compatible products, it would seem that there is wider industry acceptance for it.

2.2.2.4 Configuration Possibilities

Local area network configuration possibilities are endless and can involve heterogeneous micros, minis, word processors, and more, all within a single local area network. Telecommunications to other local area networks through gateways expand the geographic coverage of the network. Likewise, using a gateway into a wide area network provides even more possibilities. In the latter case, units in the local area network can communicate from one to another without using the resources of the wide area network and, thus, minimize the usage of mainframe resources. It seems that the key to the possibilities are the heterogeneous products which are capable of being integrated. In many cases, software has to be developed to support this, but it appears that this software is at Levels 5 through 7 of the ISO model.

2.2.3 <u>Wide Area Networks</u>

Wide area networks have been in existence for quite some time and have evolved considerably over that time. For the purposes of this discussion, a wide area network is considered to be defined as a network which is implemented using telecommunication links of the public telephone system and additional vendor-supplied hardware and software. Such a network may use microwave stations and/or satellite communications in its implementation. Given this type of implementation, the size of the geographic region which the network covers is very large. When wide area networks were first introduced, they were quite simple and allowed terminals to communicate with a single host. They have evolved over time to the point where they now are quite powerful and allow terminals, as well as hosts, to communicate with any other host in the network. Due to the evolutionary development of these network products, the vendors are limited to mainframe manufacturers with a few plug compatible offerings. Because of this, the networks tend to support only a specific manufacturers product line (or plug compatible), though the sizes of machines within the product line may range from very small to large computers. Having defined wide area networks, it is appropriate to look at the motivation for using them, their components, cite some existing products and protocols, and finally, examine the configuration possibilities.

2.2.3.1 Motivations for Using Wide Area Networks

The motivation for using wide area networks is quite simple. It is the requirement to allow terminals, hosts and other devices to communicate with geographically separated hosts. Geographical separation usually occurs in the one-mile range for the terminal to host case and much less than one mile for the host to host case.

2.2.3.2 Physical Components

The physical components associated with wide area networks are very diverse. They consist of the carrier and associated modems and devices called Interface Message Processors (IMP) in Arpanet terminology, otherwise known as a communication controller.

The carriers in wide area networks consist of telephone communication links including telephone circuits, microwave and satellite communications. Normally, telephone circuits exist at the ends of the communications link (terminal and computer) with the satellite and microwave communication somewhat transparent. These links may be either switched (e.g., dial up) or dedicated. The speed of these links generally range from 300 baud up to 9,600 baud with the higher rates requiring dedicated circuits. Through the use of multi-plexors or special lines and interface connections, higher rates of 19.2 KBPS-50 KBPS can be achieved. Depending on the implementation, more than one circuit may be required. At the end of each circuit is a modem. The modem, among other things, converts a digital signal to the analog signal required for telephone line hook up. Modems connect to terminals and communication's controllers through an RS-232C interface. For higher speed data rates (4,800 -9,600 baud), the line must be conditioned to insure that it supports the band width.

Communication's controllers may be dedicated hardware boxes or boards plugged into the CPU. Some communication's controllers are programmable and are executing code which controls the switching (message routing), while other architectures perform this function on the CPU. IBM's 3705 and CDC's 2550 are examples of dedicated communications controllers, while the DZ-11A package performs the same function on the VAX 11/780 CPU. These communication's controllers are quite generalized and can support a variety of protocols.

2.2.3.3 Products and Protocols

The protocols used by wide area networks are quite varied. Within the first three levels of the ISO model, there is a tendency for products to support standards (such as the CITT X.25 standard). However, due to the number of of users of protocols which existed prior to the development of standards, there is a tendency for the vendors to support non-standard protocols. This occurs predominantly in the second level of the ISO model. Vendors tend to support asynchronous, binary synchronous, and synchronous data link control disciplines. Beyond Level 4 in the ISO model, the vendors tend to have their own unique protocol, and furthermore, they do not always fully provide support for each level. A detailed investigation of each vendor's protocol is beyond the scope of this paper.

Three of the major mainframe vendors of hardware in the IPAD arena (CDC, DEC, and IBM) have their own set of products. Each has differing levels of capabilities as they are evolutionary products.

CDC currently provides limited wide area network capabilities through a set of network products. The major items in this set consist of the 2550/2551 network processing unit along with the Communications Control Program (CCP) and Network Access Method (NAM) product set.

DEC, through its DECNET products, provides significant wide area network capabilities. The network architecture implemented by DECNET is referred to as Digital Network Architecture. In this architecture, there are no specialized pieces of hardware front ending the host to provide network control functions as this is done by the host.

IBM, through its Advanced Communication Function product set, provides significant wide area network capabilities. The primary components in this product set are the 3704/3705 network controller and the Network Control Program (NCP)/Virtual Telecommunications Access Method (VTAM) products. Through these and related products, IBM implements its Systems Network Architecture (SNA).

A different set of products also enters the wide area network domain. These are the public/private X.25 networks marketed by different vendors. These vendors provide the first three layers of the ISO model and users link their CPUs up to them. This is an alternative to leasing long distance lines and may or may not be cost effective depending on the particular situation.

2.2.3.4 Configuration Possibilities

The configuration possibilities for wide area networks are limited by the vendor's product capabilities. Within a vendor's product line and plug compatibility, there are some degrees of freedom. Crossing vendor lines is somewhat difficult. Some specific cases are possible but these are generally limited in functionality.

Two of the more interesting possibilities involve the X.25 standard and local area networks. Depending on the capabilities of various products, it may be possible to effect an interface between heterogeneous machines using the X.25 interface standards for the first three levels of the ISO model. The remaining levels might be implemented in terms of one vendor's product set or another. Such a task is not easily accomplished, but never the less, it is appealing.

The second possible area of interest is the interfacing of a local area network to a wide area network. Such an interface would provide significant benefits for the distributed processing and distributed data base environment. This may be the more easily accomplished alternative of the two cited.

2.3 DISTRIBUTED DATA BASE MANAGEMENT SYSTEMS

As stated in section 1.2.3, distributed data base management systems have been the subject of much activity in the research community. However, commercially available production generalized distributed data base managers have yet to become a reality. Therefore, any statement as to the practicality of a particular product cannot be made.

This section will, however, examine some of the issues regarding distributed data management systems and address them from the perspective of the engineering and manufacturing environment. These issues are: the granularity of distribution; the data definition and maintenance scheme; local vs global data management; meta data maintenance; the transparency of distribution; and finally, whether or not the nature of the distribution is dynamic.

2.3.1 The Granularity of Distribution

In the distributed data base environment, subsets of the data are located at different nodes of the network. The granularity of distribution refers to how that data is divided among the nodes and what that unit of subdivision is. It is felt that the three basic alternatives for that subdivision are: the record occurrence; the record type; and the data set and/or data area. Each will be discussed.

2.3.1.1 The Record Occurrence

A record occurrence is a single retrievable unit of data. It corresponds to a row in a two-dimensional table. Using this granularity of distribution, every time a record is stored, a decision would be made by the distributed data manager as to what node in the network that record should be stored on. This decision may be based upon some value in the record, or it may be stored in the data base of the node nearest to the application causing the store. The criteria for where the record was stored would have to be defined to the DBMS. In the IPIP environment, this declaration would most likely be made in the Internal Schema if this granularity of distribution were adopted.

2.3.1.2 The Record Type

A record type is the accumulation of all records having the same definition. So if two two-dimensional tables existed, there would be two record types. The rows of one table would make up the record occurrences of one record type, and the rows of the second table would make up all the record occurrences of the second type. Again, as each individual record is stored, a decision would be made on where to store the record. However, this decision would be based upon the definition (record type) of the record occurrence. This causes all record occurrences having a like definition to be stored at a single node in the network. The distributed data manager must have the location for the storage of each record type defined to it so that it can make the appropriate decision. If this granularity of distribution were adopted, the definition to IPIP would again most likely occur in the Internal Schema.

2.3.1.3 The Data Set/Data Area

The IPAD work has defined a data set to be an arbitrary group of record occurrences from one or more record types. It represents the data associated with an engineering unit of work by task. Data sets have many properties, however, more importantly, they belong to a data area. A data area is a group of data sets or other data areas and represents the data owned by an engineering organization. In this environment, as each record is stored, the distributed data base management system would determine what data set the record belongs to and what data area that data set belonged to so that it could make a decision as to what node in the network that record would be assigned to.

In an IPIP environment, the assignment could be made in an Internal Schema declaration or via a data base administration function. In the case where no location was specified for the data set or data area, the data would be stored at the closest data base node in the network. This granularity of distribution was the one chosen for distribution during the preliminary design of IPIP. This choice was based upon many hours of discussions with the engineering staff at that time. The choice still appears to be valid in today's environment, however, this does require re-examination to determine if it is.

2.3.2 Data Maintenance Conventions

The data definition and maintenance convention is the way in which the distribution of data occurs and its integrity is assured. The three basic alternatives to be considered are: replicated data, partitioned data, and local/global data bases. The first two methods involve peer to peer relationships while the last (local/global) involves a master/subordinate relationship between the data managers. Each methodology has its advantages, disadvantages, and fits certain situations better than others. Each convention will be discussed with respect to the user data.

2.3.2.1 Replicated Data

Replicated data exists in two forms: fully replicated and partially replicated. The fully replicated data case duplicates all data values at all nodes in the network while the partially replicated case is a combination of the fully replicated case and the partitioned convention since only part of the data is replicated. This methodology provides for rapid response to all users regardless of the node they are at. The determination of what is replicated or not would be a data base administration function and would most likely be based either on definition (record type) or the data set/data area membership.

Since data is replicated at all nodes, all changes made to the contents of the data base must be reflected at all nodes. The process of maintaining the synchronization of the data bases is quite complex. Depending on requirements, synchronization could take place in real time or it could be deferred to another time, e.g., the end of the day, week, month, year, etc. Real time synchronization introduces one type of complexity while deferred synchronization introduces a different kind of complexity. Real time synchronization has an advantage in the environment where the data base must be current and available for query 99.99% of the time. If any single data base node is down due to a hardware failure, its users can be serviced by other nodes. There are some inherent disadvantages in this method, though. It generates a significant amount of network traffic and CPU activity, since every update is essentially performed on all nodes as they occur. Furthermore, the deadly embrace detection and recovery process also becomes more complex in the convention.

Deferred synchronization has an advantage where data base availability is important but the data does not have to be as current. In other words, the users may require access to data at all times but they can work with data which is not totally up to date. Synchronization in this environment may occur at a low activity time, so the increased CPU activity may not be as important as it is with real time synchronization. The actual synchronization of the data base becomes quite complex, however, and may be data and application dependent as well as time dependent. Furthermore, redundant updates may occur as in the case where the same error is corrected at multiple nodes.

A disadvantage associated with replicated data is the required disk storage. Regardless of how much data is replicated, some disk storage contains the same data.

In general, it is not well understood if and where the replicated form of distributed data fits in the engineering and manufacturing environment. Moreover, if it fits the environment, the synchronization requirements are also not well understood. These areas need more investigation.

2.3.2.2 Partitioned Data

The partitioned data convention duplicates no user data at the data base nodes of the network but rather routes requests for data to the node in the network where the data is resident. The location of the data in the network could come either directly from the user or the distributed data manager directories based upon previous user or data administrator input. Having the data manager provide locations of the data would be preferable. This particular convention seems to fit the required environment much better since it avoids the problems of replicated data especially with respect to storage requirements. Furthermore, if the data set/data area is the granularity of distribution, then the number of multi-node data base requests would most likely be reduced since the data could be located on nodes in the network based upon access patterns as well as the location of user engineers and applications.

2.3.2.3 The Local/Global Data Base Management System

This methodology has been alluded to before in section 1.2.3 and may be considered as a combination of replicated and partitioned data. Using this methodology, the global data base would be designated as the official data base. The global data base could be distributed itself by using one of the two methods just cited. Users would extract part of the data base and move it to a local data manager which is resident on a local CPU. The CPU might be a micro-computer driven engineering work station. The user would make the changes necessary on the local data base, and then at some later point, the local and global data bases would be reconciled. The data set is the user's natural unit of data to down load to the local data manager since it is task-oriented. In this environment, the user would control all distribution activity between the local and global data managers. It is also conceivable that the data set would be passed on to several other local hosts prior to being reconciled with the global data manager. A critical issue in this convention is the prevention of multiple users who have been given update permission from "checking out" the same data set making separate updates and then overlaying each other's changes. The update privilege is granted by owners of the data, but the data manager must enforce the owner's permissions.

The combination of local/global data management techniques along with partitioned data between peer distributed data managers seems to provide a good solution to the integrated engineering and manufacturing environment if the data set/data area is the granularity of distribution. This appears to be especially true for data which represents or deals with manufactured items even if it is obtained from a design drafting system or is input to NC/DNC machines.

2.3.3 Meta Data Maintenance

Meta data is the data about data. It includes both system meta data and user meta data. The maintenance of these also becomes complicated in the distributed data base environment.

2.3.3.1 System Meta Data

System meta data is data which the DBMS creates about the data it stores. It includes such items as directory information and definition information. Definition information, such as item and record definitions, does not necessarily change in content as a result of distribution. Directory information which contains things like location of data on the disk may change in content as a result of distribution. Depending on the implementation, distributed data may require peer data managers to contain the node location of all data in the system. Another implementation may broadcast requests to all data managers, and this would require each data manager to know what its contents are. Since both types of meta data are data to the system, things like the granularity of distribution and the data maintenance convention must be examined for these items also. Preliminary analysis of this issue has suggested that some system meta data must be fully replicated. This issue requires much further analysis and design before a final method is chosen.

2.3.3.2 User Meta Data

User meta data contains user assigned information about data. This information is most logically associated with the granuals of distribution. An example of this might be the "released status" of a data set. This meta data could be either fully replicated or partitioned. If the user meta data is partitioned, the meta data data for a particular data set might be located on the same node that data set is located on. Which option is chosen is dependent upon how often the data is updated in contrast to how often it is queried. More analysis is required in this area as well, although it appears that the granularity of distribution is at the record occurrence for user meta data.

2.3.4 Transparency of Distribution

The transparency of distribution refers to whether or not the user is required to know the location of the data. The appropriate direction in this area appears to be based upon a combination of transparent and non-transparent distribution. The distribution should be transparent between peer global data managers. However, it should not be transparent between the local and global data managers, requiring users to take specific actions to move data between global and local data bases.

Having transparent distribution between peer global data managers does not imply that the data moves from one node to another node in the network based upon usage. It should always take a specific action to change the location of data. For the set of peer data managers which comprise global data bases, that action should be a data base administration function.

2.3.5 The Data Base Administration Function

In the distributed data base environment, the data base administration function should also be distributed. The distribution of this function should correspond to organizational lines so that managers of groups can have local control of the data base administration function just as they have control over the local computer.

2.4 OTHER SOFTWARE

To this point, software involving both networks and data base management systems has been discussed. There are two other major items of software involved in distributed environments.

2.4.1 Teleprocessing Monitors

Teleprocessing (TP) monitors facilitate the development of certain types of online production systems. This environment has been called transaction processing in previous IPAD work. These products support the environment where several people are executing the same program in a highly repetitive fashion. The advantages that these products offer is that they generally allow multiple users to share the same code but the variables of each user are separated. This makes more effective use of the core. Additionally, they provide subtasking capabilities so that users share control points or regions of memory thus reducing the number of regions or control points required by the total online user community. Common examples of these and similar products are CDC's TAFS and IBM's CICS. These products may be used in the integrated engineering and manufacturing environment.

2.4.2 Operating Systems

Just as the hardware in the engineering and manufacturing environment is heterogeneous, so are the operating systems. Some vendors offer more than one operating system within a given product line introducing additional levels of complication. The software (such as the data base management system) used on the various operating systems should exploit the inherent capability of that operating system but should preserve similiar functionality and interfaces where possible.

3.0 THE ROLE OF WORK STATION IN THE DISTRIBUTED ENGINEERING ENVIRONMENT

The last few years have seen an explosion in microcomputer technology. The micro has evolved from an 8-bit toy to a 32-bit workhorse. This evolution has opened the door for many uses of micro-based work stations in the engineering and manufacturing environment. The same time frame has seen the development of small 32-bit mini's (such as the VAX 11/730). This development has opened the door even further for local engineering work stations.

3.1 TYPES OF WORK STATIONS

In general, there are three types of work stations in the engineering environment: the design drafting system, office automation stations, and the general purpose micro-based work station. Each provides unique information processing capabilities. The ultimate goal is to integrate these capabilities.

3.1.1 The Design Drafting System

The design drafting system is generally a self-contained turnkey system with integrated hardware and software. These systems are used in the design of parts and the creation of part geometry. Some of them can be used to generate NC tapes to control the manufacture of those parts. In most cases, these systems contain their own specialized data structures and geometry representations. Therefore, the integration of these devices into other hardware and software systems proves to be somewhat difficult.

Another entry in the market is the specialized intelligent graphics display work station such as the SEILLAC. These devices provide good capabilities for the creation and display of engineering geometry. These graphics display work stations also seem to have the capability to emulate dumb terminals, however, they do not appear to be able to act as an intelligent node in a network.

3.1.2 Office Automation Stations

Office automation stations generally can be classified as either word processors or professional work stations. Discussion of these is included because the engineering environment not only includes specific design and analysis processes but also includes the communication and documentation aspects of those processes to other engineering organizations as well as organizational management.

Word processors (such as the XEROX 860) were designed to perform general word processing functions such as document creation, editing, and storage. As these devices are maturing, their capabilities are improving. Depending on the particular brand and model, it may be possible to get a widely used operating system for them, e.g., CP/M, MS/DOS. This makes available to them large numbers of existing microcomputer programs. Again, depending on the model, it may be possible to obtain specialized application software for them such as electronic spread sheets. These added capabilities open up whole new arenas for their use.

The professional work station such as the XEROX Star and Lisa by Apple provide broader capabilities. In addition to word processing functions, they may provide electronic mail functions as well as chart and view foil preparation capabilities. Depending on the printer supporting them, print font size and style selection may also be provided.

3.1.3 General Purpose Microcomputer

The general purpose microcomputer is the area which has had the most evolution in recent years. The success of Apple and Radio Shack have drawn in mainframe manufacturers like CDC, DEC, and IBM. Microcomputers started out as 8-bit machines using one of three basic CPU chips. The CP/M operating system proved to become a standard operating system creating an environment of portable software. These micros are single-user systems and provide a variety of capabilities depending on the software packages installed, such as word processing, electronic spread sheets, graphics, and many more specific applications. Depending on the vendor's marketing strategy, there may be a wide variety of hardware from plug compatible vendors. This hardware generally expands or complements the hardware provided by the original vendor. Apple and IBM are the two vendors which exemplify this approach while Radio Shack is the opposite side of the continuum. Apple and IBM invite plug-compatible devices while Radio Shack discourages them.

Late 1981 and early 1982 saw the introduction of the 16-bit micro from companies like Radio Shack, IBM, and DEC. Around the same time, the 5¼ inch hard disk based on Winchester technology was introduced. These two advancements have significantly increased the capability of these machines. Whereas most 8-bit machines had address space limitations of 64 K bytes, the 16-bit machines have a theoretical address space of 1 million bytes or larger. This allows the use of larger programs. The hard disk provides faster data access and thus improves response time.

Operating systems for micros are advancing similarly with the development of multi-tasking operating systems (such as CP/M 86). These operating systems allow a single user to do one or more background activities such as printing of a program, compiling another, while the user is editing a data file. A capability offered by other operating systems is the support of multiple users from a single micro.

The 32-bit micro is just around the corner. Some are available today (such as the HP 9000). In general, it won't be long before 32-bit micros have reached a state of maturity.

In the engineering environment, micros can be utilized for such tasks as data definition, collection, validation, display, modification, and other such data management activities. By moving these activities off the host to micros allows the large host to be more effectively matched to the tasks.

3.2 CONFIGURATION POSSIBILITIES

The configuration possibilities involving work stations are widely varied and provide for a large number of combinations. There are, however, three basic types of configurations: the standalone work station; the work station in a local area network; and the work station in a wide area network.

3.2.1 Standalone Work Station

The standalone work station occurs in two varieties, either single user or multi-user. In either situation, the work station must be somewhat self-contained. The only type of communication between it and any other type of computer (micro, mini, or mainframe) would be in the form of a file transfer for data or as a dumb terminal for interaction with the other system. This configuration, although simple to employ, does not fully exploit the capabilities of the work station. In most cases, the communication link would be a dial up line. Transfer of data between work stations in this manner is awkward, particularly if there are two work stations involved.

3.2.2 Work Stations in Local Area Networks

In this configuration, work stations would be attached to a local area network via the appropriate NIU's and taps. This configuration would provide a basis for communication between heterogeneous work stations without dial up lines and also provide for the sharing of certain kinds of peripheral devices such as printers. Furthermore, the work station may also communicate with hosts which are also part of the local area network. If communications are desired with work stations on other local area networks or hosts not on the local area network, gateways can be employed. A gateway is an NIU which takes network signals and routes them over telephone lines. It appears that this configuration is most beneficial in an environment where a large number of work stations are in close physical proximity to one another and require large amounts of communication between each other.

3.2.3 Work Stations in Wide Area Networks

In this configuration, the work stations would be defined to the wide area network as a node in the network. In such an environment, the capabilities are very much limited to the support that the vendor's wide area network provides. Communications would primarily be between the work station and the supporting node. If communications were required between other work stations, they would have to be routed to the other work station by the supporting node.

This configuration tends to be best suited for environments where the work stations are geographically dispersed and communications are primarily between the work station and its supporting host/node.

3.3.4 Network Combinations

Since neither local nor wide area networks alone may satisfy a firm's requirements, combinations of the two may be necessary. This would support physically concentrated and dispersed work stations. Such a combination may be possible, but the extent by which it may be accomplished is a function of the vendor-provided software and limitations which may exist.

4.0 POSSIBLE ARCHITECTURE FOR DISTRIBUTED COMPUTING

Recall figure 1.3-1 in which the concept of a layered network was depicted. This section will present a possible architecture for distributed computing evolving from the current IPAD environment. This architecture will utilize existing IPAD products and the items discussed in sections 2.0 and 3.0. Figure 4.0-1 depicts this architecture which is an expansion of the current architecture shown in figure 4.0-2. It includes heterogeneous mainframes, multiple wide and local area networks. The CDC CYBER computer and the IBM represent the computers at the center of the concentric circles. The VAX 11/780 represents the size of the computer at second level and finally the work stations represent the outer layer of hardware. The hardware configuration far exceeds the required resources for IPAD development. Furthermore, this configuration does not fit any specific situation. It is present for concept development purposes and contains elements which may be applied to specific problem areas in any firm. The following sections define a phased approach for moving toward that architecture. Each phase will be discussed in terms of architecture and functionality.

4.1 PHASE I

Phase I provides work in wide area networking and distributed data base management utilizing a local/global data base management concept.

4.1.1 Architecture

The first phase shown in figure 4.1-1 adds general purpose micro-based work stations to the existing environment. These work stations would be connected via telecommunication links to the VAX 11/780. They would possibly be defined as nodes in a DECNET network. This would make them nodes in a wide area network as discussed.

IPIP would be resident on the CYBER and represent a global data manager. A local data manager would be resident on the work station or possibly the VAX.

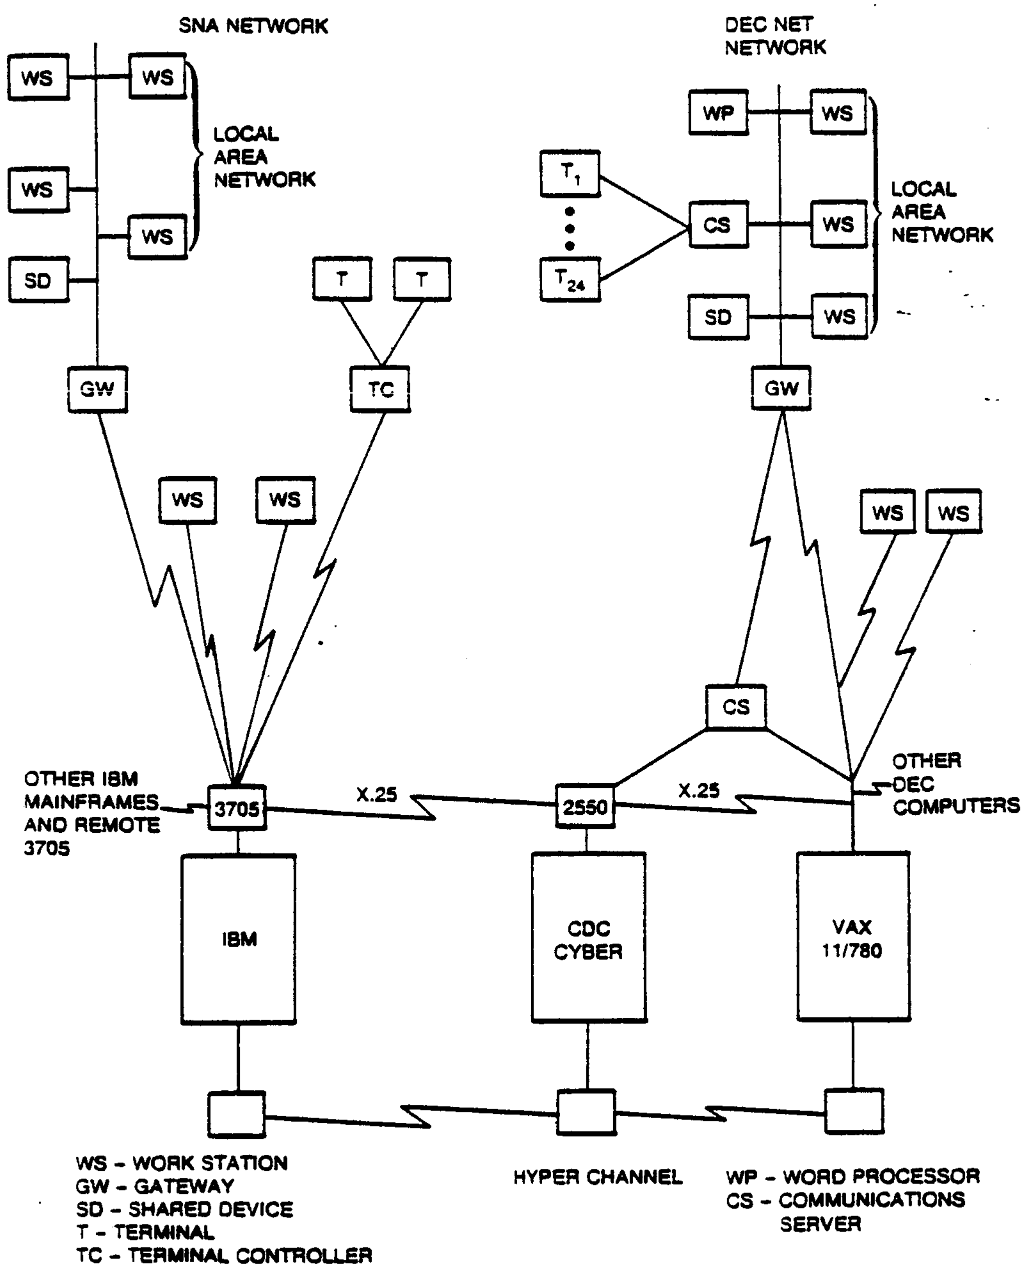

FIGURE 4.0-1 PHASE III ARCHITECTURE

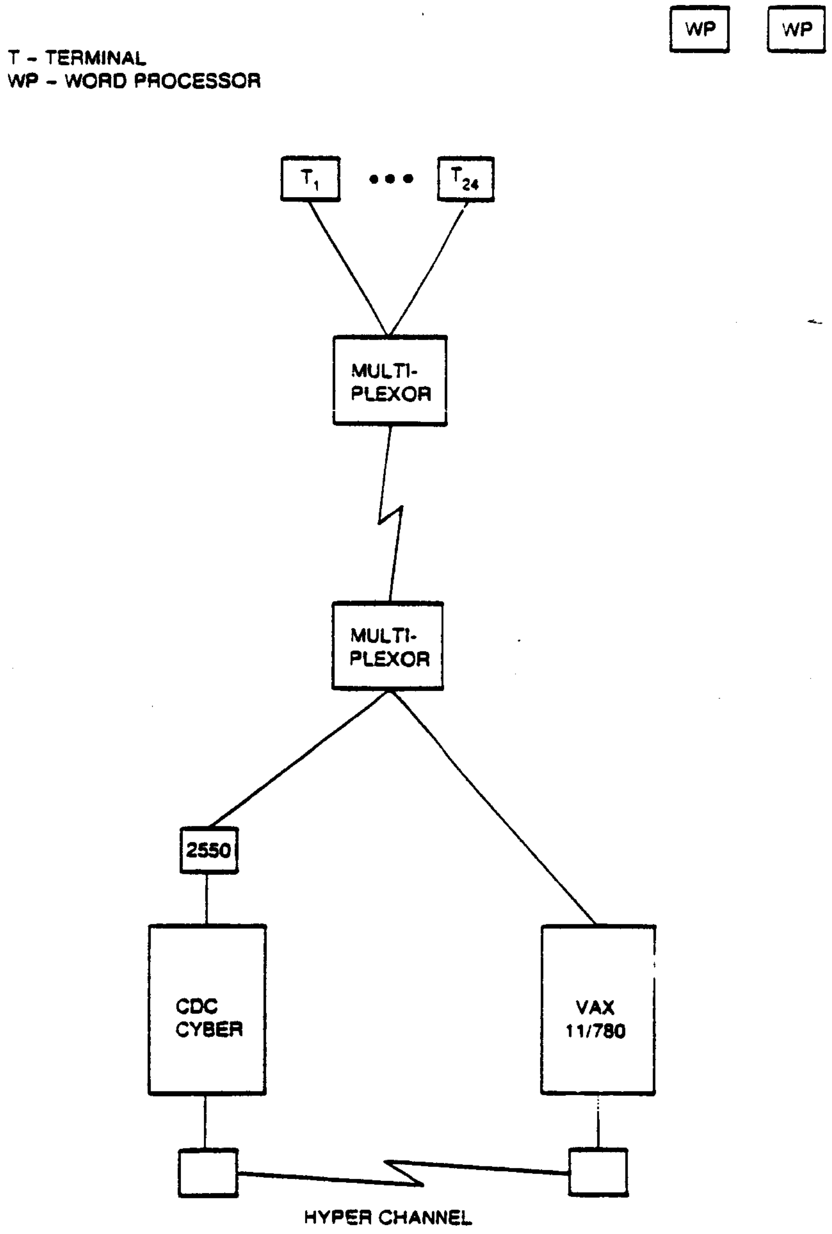

FIGURE 4.0-2 CURRENT ARCHITECTURE

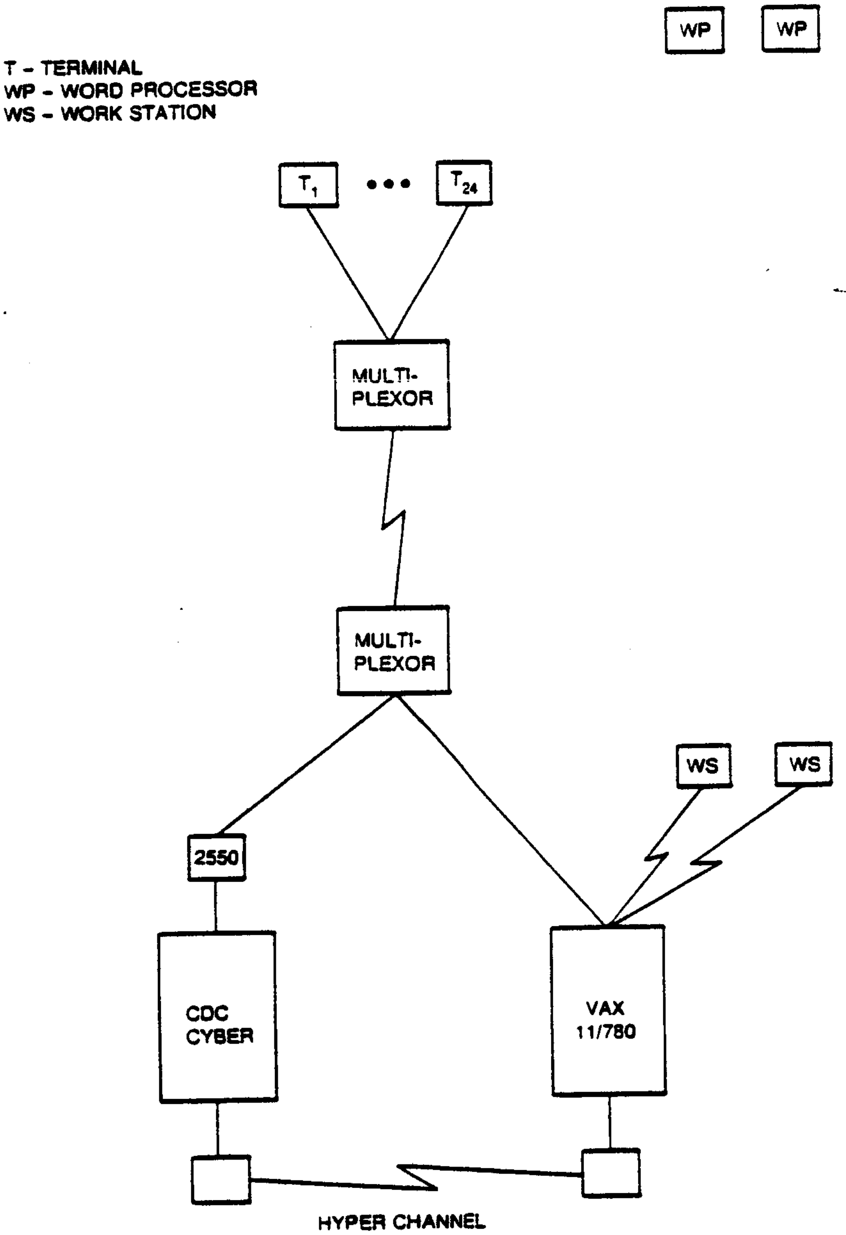

FIGURE 4.1-1 PHASE I ARCHITECTURE

4.1.2 Functionality

The user would be able to request that a data set (the granularity of distribution) be down loaded from the global data manager to the local data manager. The data would be transferred from the CYBER via the VAX 11/780 to the work station. Once on the work station, the user would be able to modify it as desired. When the modification was complete, it would be returned to the global data manager.

Also in the environment, would be the support of distributed functions, allowing an application program to execute on the work station, and request data from the global data manager.

4.2 PHASE II

Phase II introduces a local area network to the environment established in Phase I and explores the interconnection of local area networks and wide area networks. It would also begin work on peer distributed data managers and the use of transaction processing. Figure 4.2-1 illustrates the network after Phase II.

4.2.1 Architecture

The architecture would evolve in two distinct segments during this phase. The first segment would introduce a local area network to provide communications between micro-based work stations and word processors. They would also be able to share a printer. A DECNET gateway would be utilized to provide communications to the VAX 11/780. Once this architecture was established and performance of the network understood, the existing ASCII terminals would be connected to the local area network via communication servers. In addition to this, a communication server would be utilized to front end the CYBER and VAX. This would allow replacement of the existing multi-plexors.

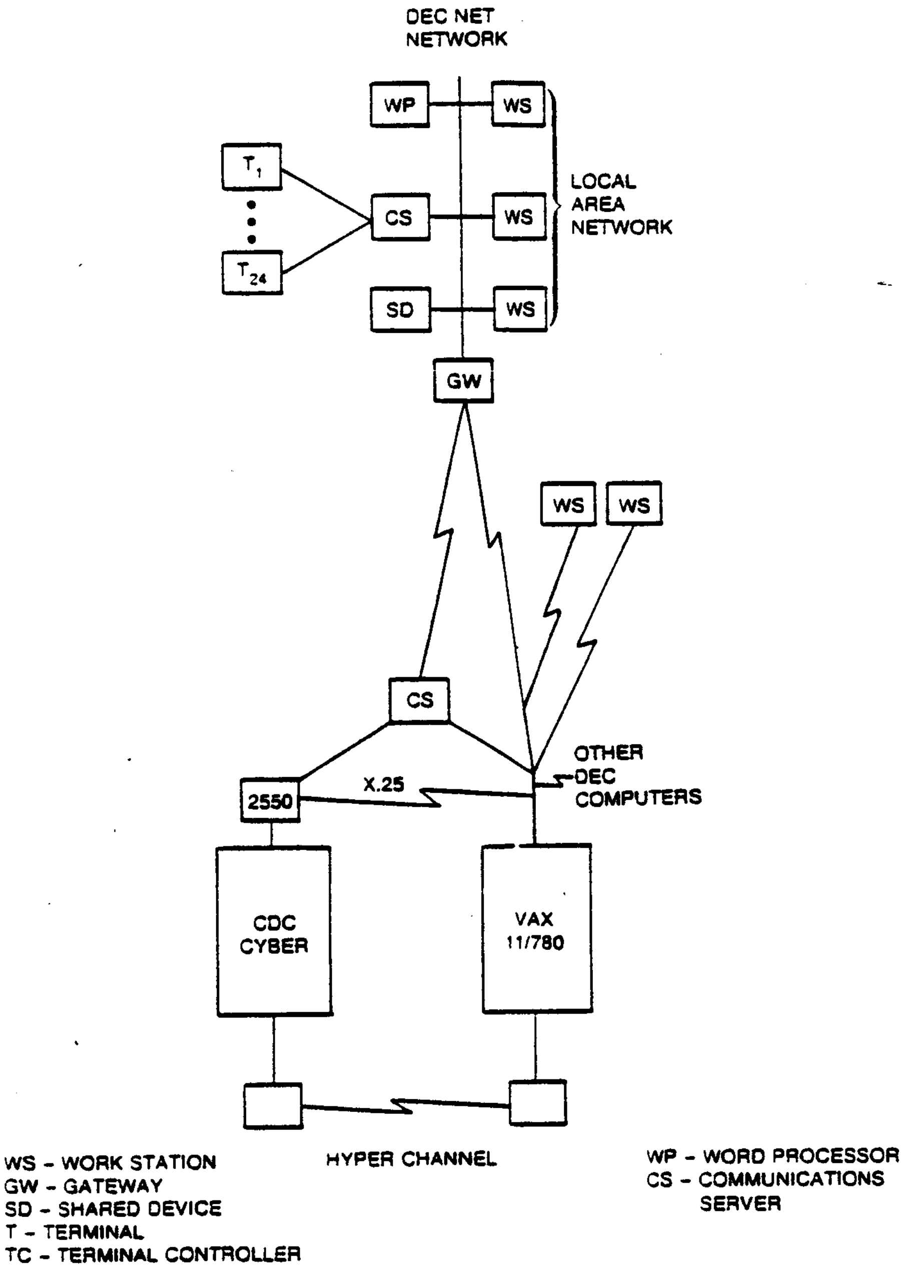

FIGURE 4.2-1 PHASE II ARCHITECTURE

4.2.2 Functionality

In this environment, information would be able to be exchanged between all work stations. The micro-based work stations would support the local data managers and be able to interact with the global data manager. Once the second segment was completed, the ASCII terminals would be able to access either the VAX or CYBER via the local area network. Also, the work stations would gain access to the CYBER in a direct manner through the local area network. The shared printer should be able to support the printing of output from either the VAX, CYBER, or any of the work stations. The peer to peer data managers would be able to route requests to other peer data managers and respond to their requests, thus, allowing multiple peers to respond to a single request from local data managers.

Transaction processing would allow for typical online transactions and may be based on the CDC product TAFS. This may require more intelligent terminals.

4.3 PHASE III

Phase III would introduce an IBM host, explore the interconnection of wide area networks via X.25 interfaces, further explore distributed data management, and transaction processing. This brings the architecture to the level shown in figure 4.0-1.

4.3.1 Architecture

The two major networking facets of this phase are the introduction of an IBM SNA network which supports a second local area network. If this local area network was the same as in Phase II, then communications could be established between the two. Also of importance is the X.25 links among the IBM, CDC, and VAX hosts. Continuing with the local area network philosophy for connection of heterogeneous hosts, the hyperchannel network would extend to the IBM. The peer distributed data manager would move to the IBM and/or VAX node and at that point, peer data managers would possibly exist on

CYBER, IBM, and VAX hosts.

4.3.2 Functionality

In this final environment, any node would be able to request data from any other nodes containing the peer data managers. Any local data manager node would be able to transfer data to any other local data manager node. Communications from one work station to another could be between two local area networks via telecommunication link, X.25 link, or hyperchannel.

4.4 CONCLUSION

The architecture expressed here certainly achieves the goals of distributed data base management between heterogeneous machines. Each phase described is a logical extension of the previous phase. However, the phases are quite ambitious and will require a significant amount of effort. The investment may be high, but the returns would also be very high.